


[> home](#) [> about](#) [> feedback](#) [> login](#)

US Patent &amp; Trademark Office



Try the *new* Portal design  
Give us your opinion after using it.

## Search Results

Search Results for: [scoreboard and electronic\* ]  
Found 46 of 126,269 searched.

Search within Results

[> Advanced Search](#)  
[> Search Help/Tips](#)

Sort by: [Title](#) [Publication](#) [Publication Date](#) [Score](#)

Results 1 - 20 of 46 [short listing](#)

1 2 3

- 09/25/2007
- 1 TSEM, a flexible scenario based small forces model 97%  
☒ Bruce D. Link , Henry D. Shapiro  
**Proceedings of the 11th conference on Winter simulation - Volume 2** December 1979  
 An accurate computer model of a small force engagement is useful in evaluating the combat effectiveness of armed escorts for sensitive shipments, security guard forces and military patrols. The Transportation Safeguards Effectiveness Model (TSEM), primarily intended for, but not limited to, the study of ambushes of armed convoys, provides the user with considerably greater flexibility in directing the actions of the combatants than previous models. A user oriented script language is present ...
  - 2 The Psychology of How Novices Learn Computer Programming 95%  
☒ Richard E. Mayer  
**ACM Computing Surveys (CSUR)** January 1981  
 Volume 13 Issue 1
  - 3 Talisman: fast and accurate multicomputer simulation 89%  
☒ Robert C. Bedichek  
**ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1995 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems** May 1995  
 Volume 23 Issue 1  
 Talisman is a simulator that models the execution semantics and timing of a multicomputer. Talisman is unique in combining high semantic accuracy, high timing accuracy, portability, and good performance. This good performance allows users to run significant programs on large simulated multicomputers. The combination of high accuracy and good performance yields an ideal tool for evaluating architectural trade-offs. Talisman models the semantics of virtual memory, a circuit-switched interno ...
  - 4 Gun turret drive stabilization system for the M1A2 88%  
☒ Richard E. Copra , Joseph R. Vogl  
**Proceedings of the conference on TRI-Ada '91: today's accomplishments; tomorrow's**

## expectations December 1991

- 5 Exploiting fine-grain thread level parallelism on the MIT multi-ALU processor 87%  
[A] Stephen W. Keckler , William J. Dally , Daniel Maskit , Nicholas P. Carter , Andrew Chang , Whay S. Lee  
**ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture April 1998**  
Volume 26 Issue 3  
Much of the improvement in computer performance over the last twenty years has come from faster transistors and architectural advances that increase parallelism. Historically, parallelism has been exploited either at the instruction level with a grain-size of a single instruction or by partitioning applications into coarse threads with grain-sizes of thousands of instructions. Fine-grain threads fill the parallelism gap between these extremes by enabling tasks with run lengths as small as 20 cyc ...
- 6 Providing location information in a ubiquitous computing environment (panel session) 85%  
[A] Mike Spreitzer , Marvin Theimer  
**ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM symposium on Operating systems principles December 1993**  
Volume 27 Issue 5  
To take full advantage of the promise of ubiquitous computing requires the use of location information, yet people should have control over who may know their whereabouts. We present an architecture that achieves these goals for an interesting set of applications. Personal information is managed by User Agents, and a partially decentralized Location Query Service is used to facilitate location-based operations. This architecture gives users primary control over their location information, at the ...
- 7 Efficient dynamic scheduling through tag elimination 84%  
[A] Dan Ernst , Todd Austin  
**ACM SIGARCH Computer Architecture News May 2002**  
Volume 30 Issue 2  
An increasingly large portion of scheduler latency is derived from the monolithic content addressable memory (CAM) arrays accessed during instruction wakeup. The performance of the scheduler can be improved by decreasing the number of tag comparisons necessary to schedule instructions. Using detailed simulation-based analyses, we find that most instructions enter the window with at least one of their input operands already available. By putting these instructions into specialized windows with fe ...
- 8 The M-Machine multicomputer 84%  
[A] Marco Fillo , Stephen W. Keckler , William J. Dally , Nicholas P. Carter , Andrew Chang , Yevgeny Gurevich , Whay S. Lee  
**Proceedings of the 28th annual international symposium on Microarchitecture December 1995**
- 9 A Conceptual Framework for Computer Architecture 83%  
[A] S. S. Reddi , E. A. Feustel  
**ACM Computing Surveys (CSUR) June 1976**  
Volume 8 Issue 2
- 10 Psychology of calculator languages: a framework for describing differences in users' knowledge 82%  
[A] Richard E. Mayer , Piraye Bayman  
**Communications of the ACM August 1981**  
Volume 24 Issue 8

This paper presents a framework for describing users' knowledge of how a simple four-function calculator operates. Differences among novices and experts in their conceptions of "what goes on inside the calculator" for various sequences of button presses are summarized. Individual differences include different views on when an expression is evaluated, different procedures for evaluating a chain of arithmetic, and different rules for evaluating unusual sequences of key presses.

- 11 An overview of airland combat modeling and simulation 82%  
☐ Michael W. Garrambone  
**Proceedings of the 24th conference on Winter simulation** December 1992
  
- 12 Using object modeling to transform structured analysis into object oriented design 82%  
☐ G. K. Khalsa  
**Proceedings of the sixth Washington Ada symposium on Ada** July 1989
  
- 13 Tutorial: Compiling concurrent languages for sequential processors 80%  
☐ Stephen A. Edwards  
**ACM Transactions on Design Automation of Electronic Systems (TODAES)** April 2003  
 Volume 8 Issue 2  
 Embedded systems often include a traditional processor capable of executing sequential code, but both control and data-dominated tasks are often more naturally expressed using one of the many domain-specific concurrent specification languages. This article surveys a variety of techniques for translating these concurrent specifications into sequential code. The techniques address compiling a wide variety of languages, ranging from dataflow to Petri nets. Each uses a different method, to some degr ...
  
- 14 Session 2: embedded system techniques (1): A case for dynamic pipeline scaling 80%  
☐ Jinson Koppanalil , Prakash Ramrakhyani , Sameer Desai , Anu Vaidyanathan , Eric Rotenberg  
**Proceedings of the international conference on Compilers, architecture, and synthesis for embedded systems** October 2002  
 Energy consumption can be reduced by scaling down frequency when peak performance is not needed. A lower frequency permits slower circuits, and hence a lower supply voltage. Energy reduction comes from voltage reduction, a technique called Dynamic Voltage Scaling (DVS). This paper makes the case that the useful frequency range of DVS is limited because there is a lower bound on voltage. Lowering frequency permits voltage reduction until the lowest voltage is reached. Beyond that point, lowering ...
  
- 15 Improving TCP performance over mobile networks 80%  
☐ Hala Elaarag  
**ACM Computing Surveys (CSUR)** September 2002  
 Volume 34 Issue 3  
*Transmission Control Protocol (TCP)* is the most commonly used transport protocol on the Internet. All indications assure that mobile computers and their wireless communication links will be an integral part of the future internetworks. In this paper, we present how regular TCP is well tuned to react to packet loss in wired networks. We then define mobility and the problems associated with it. We discuss why regular TCP is not suitable for mobile hosts and their wireless links by p ...
  
- 16 Illustrative risks to the public in the use of computer systems and related technology 80%  
☐ Peter G. Neumann  
**ACM SIGSOFT Software Engineering Notes** January 1996  
 Volume 21 Issue 1

- 17 Inhabited television: broadcasting interaction from within collaborative virtual environments 80%  
[A] ACM Transactions on Computer-Human Interaction (TOCHI) December 2000  
Volume 7 Issue 4  
Inhabited television combines collaborative virtual environments (CVEs) with broadcast television so that on-line audiences can participate in television shows within shared virtual worlds. We describe a series of experiments with inhabited television, beginning with the NOWninty6 poetry performance, The Mirror, and Heaven & Hell&mdash;Live. These early experiments raised fundamental questions for inhabited television concerning the extent to which it is possible to establish fast-pace ...
- 18 Associative and Parallel Processors 80%  
[A] Kenneth J. Thurber , Leon D. Wald  
ACM Computing Surveys (CSUR) December 1975  
Volume 7 Issue 4
- 19 Instruction set selection for ASIP design 80%  
[A] Michael Gschwind  
Proceedings of the seventh international workshop on Hardware/software codesign March 1999
- 20 On the scheduling of variable latency functional units 80%  
[A] Silva M. Mueller  
Proceedings of the eleventh annual ACM symposium on Parallel algorithms and architectures June 1999

---

Results 1 - 20 of 46    short listing

◀  
Prev  
Page

1   2   3

▶  
Next  
Page

---

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.


[> home](#) [> about](#) [> feedback](#) [> login](#)

US Patent &amp; Trademark Office



Try the *new* Portal design  
Give us your opinion after using it.


## Search Results

Search Results for: **[(updat\* or modify\* delet\* or insert\* or version\*)<AND>((report\*<AND>((scoreboard and electronic\* ) ) ) )]**  
Found 29 of 126,269 searched.

## Search within Results



[> Advanced Search](#)[> Search Help/Tips](#)





---

Sort by: [Title](#) [Publication](#) [Publication Date](#) [Score](#)  [Binder](#)


---

Results 1 - 20 of 29    [short listing](#)

 [Prev Page](#)    1    2    [Next Page](#) 

- 
- |       |  |     |
|-------|--|-----|
| 1     | <a href="#">TSEM, a flexible scenario based small forces model</a><br> Bruce D. Link , Henry D. Shapiro<br><b>Proceedings of the 11th conference on Winter simulation - Volume 2</b> December 1979<br>An accurate computer model of a small force engagement is useful in evaluating the combat effectiveness of armed escorts for sensitive shipments, security guard forces and military patrols. The Transportation Safeguards Effectiveness Model (TSEM), primarily intended for, but not limited to, the study of ambushes of armed convoys, provides the user with considerably greater flexibility in directing the actions of the combatants than previous models. A user oriented script language is present ...   | 97% |
| <hr/> |  |     |
| 2     | <a href="#">The Psychology of How Novices Learn Computer Programming</a><br> Richard E. Mayer<br><b>ACM Computing Surveys (CSUR)</b> January 1981<br>Volume 13 Issue 1  | 95% |
| <hr/> |  |     |
| 3     | <a href="#">Talisman: fast and accurate multicomputer simulation</a><br> Robert C. Bedichek<br><b>ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1995 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems</b> May 1995<br>Volume 23 Issue 1<br>Talisman is a simulator that models the execution semantics and timing of a multicomputer. Talisman is unique in combining high semantic accuracy, high timing accuracy, portability, and good performance. This good performance allows users to run significant programs on large simulated multicomputers. The combination of high accuracy and good performance yields an ideal tool for evaluating architectural trade-offs. Talisman models the semantics of virtual memory, a circuit-switched interno ... | 89% |
| <hr/> |  |     |
| 4     | <a href="#">Exploiting fine-grain thread level parallelism on the MIT multi-ALU processor</a>  | 87% |

- 4 Stephen W. Keckler , William J. Dally , Daniel Maskit , Nicholas P. Carter , Andrew Chang , Whay S. Lee  
**ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture** April 1998  
Volume 26 Issue 3  
Much of the improvement in computer performance over the last twenty years has come from faster transistors and architectural advances that increase parallelism. Historically, parallelism has been exploited either at the instruction level with a grain-size of a single instruction or by partitioning applications into coarse threads with grain-sizes of thousands of instructions. Fine-grain threads fill the parallelism gap between these extremes by enabling tasks with run lengths as small as 20 cyc ...
- 5 Providing location information in a ubiquitous computing environment (panel session) 85%  
4 Mike Spreitzer , Marvin Theimer  
**ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM symposium on Operating systems principles** December 1993  
Volume 27 Issue 5  
To take full advantage of the promise of ubiquitous computing requires the use of location information, yet people should have control over who may know their whereabouts. We present an architecture that achieves these goals for an interesting set of applications. Personal information is managed by User Agents, and a partially decentralized Location Query Service is used to facilitate location-based operations. This architecture gives users primary control over their location information, at the ...
- 6 Efficient dynamic scheduling through tag elimination 84%  
4 Dan Ernst , Todd Austin  
**ACM SIGARCH Computer Architecture News** May 2002  
Volume 30 Issue 2  
An increasingly large portion of scheduler latency is derived from the monolithic content addressable memory (CAM) arrays accessed during instruction wakeup. The performance of the scheduler can be improved by decreasing the number of tag comparisons necessary to schedule instructions. Using detailed simulation-based analyses, we find that most instructions enter the window with at least one of their input operands already available. By putting these instructions into specialized windows with fe ...
- 7 The M-Machine multicomputer 84%  
4 Marco Fillo , Stephen W. Keckler , William J. Dally , Nicholas P. Carter , Andrew Chang , Yevgeny Gurevich , Whay S. Lee  
**Proceedings of the 28th annual international symposium on Microarchitecture** December 1995
- 8 Psychology of calculator languages: a framework for describing differences in users' knowledge 82%  
4 Richard E. Mayer , Piraye Bayman  
**Communications of the ACM** August 1981  
Volume 24 Issue 8  
This paper presents a framework for describing users' knowledge of how a simple four-function calculator operates. Differences among novices and experts in their conceptions of "what goes on inside the calculator" for various sequences of button presses are summarized. Individual differences include different views on when an expression is evaluated, different procedures for evaluating a chain of arithmetic, and different rules for evaluating unusual sequences of key presses.
- 9 An overview of airland combat modeling and simulation 82%

- Michael W. Garrambone  
**Proceedings of the 24th conference on Winter simulation** December 1992
- 10 Tutorial: Compiling concurrent languages for sequential processors 80%  
Stephen A. Edwards  
**ACM Transactions on Design Automation of Electronic Systems (TODAES)** April 2003  
Volume 8 Issue 2  
Embedded systems often include a traditional processor capable of executing sequential code, but both control and data-dominated tasks are often more naturally expressed using one of the many domain-specific concurrent specification languages. This article surveys a variety of techniques for translating these concurrent specifications into sequential code. The techniques address compiling a wide variety of languages, ranging from dataflow to Petri nets. Each uses a different method, to some degree ...
- 11 Session 2: embedded system techniques (1): A case for dynamic pipeline scaling 80%  
Jinson Koppanalil, Prakash Ramrakhyani, Sameer Desai, Anu Vaidyanathan, Eric Rotenberg  
**Proceedings of the international conference on Compilers, architecture, and synthesis for embedded systems** October 2002  
Energy consumption can be reduced by scaling down frequency when peak performance is not needed. A lower frequency permits slower circuits, and hence a lower supply voltage. Energy reduction comes from voltage reduction, a technique called Dynamic Voltage Scaling (DVS). This paper makes the case that the useful frequency range of DVS is limited because there is a lower bound on voltage. Lowering frequency permits voltage reduction until the lowest voltage is reached. Beyond that point, lowering ...
- 12 Improving TCP performance over mobile networks 80%  
Hala Elaarag  
**ACM Computing Surveys (CSUR)** September 2002  
Volume 34 Issue 3  
*Transmission Control Protocol (TCP)* is the most commonly used transport protocol on the Internet. All indications assure that mobile computers and their wireless communication links will be an integral part of the future internetworks. In this paper, we present how regular TCP is well tuned to react to packet loss in wired networks. We then define mobility and the problems associated with it. We discuss why regular TCP is not suitable for mobile hosts and their wireless links by p ...
- 13 Illustrative risks to the public in the use of computer systems and related technology 80%  
Peter G. Neumann  
**ACM SIGSOFT Software Engineering Notes** January 1996  
Volume 21 Issue 1
- 14 Inhabited television: broadcasting interaction from within collaborative virtual environments 80%  
ACM Transactions on Computer-Human Interaction (TOCHI) December 2000  
Volume 7 Issue 4  
Inhabited television combines collaborative virtual environments (CVEs) with broadcast television so that on-line audiences can participate in television shows within shared virtual worlds. We describe a series of experiments with inhabited television, beginning with the NOWninetysix poetry performance, The Mirror, and Heaven & Hell&mdash;Live. These early experiments raised fundamental questions for inhabited television concerning the extent to which it is possible to establish fast-pace ...
- 15 Associative and Parallel Processors 80%  
Kenneth J. Thurber, Leon D. Wald  
**ACM Computing Surveys (CSUR)** December 1975

## Volume 7 Issue 4

- 16 On the scheduling of variable latency functional units 80%  
[A] Silva M. Mueller  
**Proceedings of the eleventh annual ACM symposium on Parallel algorithms and architectures** June 1999
- 17 Introduction to ProcessModel and ProcessModel 9000 80%  
[A] John F. Vick , Charles R. Harrell  
**Proceedings of the 30th conference on Winter simulation** December 1998
- 18 Introduction to ProcessModel and ProcessModel 9000 80%  
[A] Bruce D. Gladwin , Charles R. Harrell  
**Proceedings of the 29th conference on Winter simulation** December 1997
- 19 The computer bowl 77%  
[A] Karen A. Frenkel  
**Communications of the ACM** February 1989  
Volume 32 Issue 1
- 20 SALSA: the stochastic approach for link-structure analysis 77%  
[A] R. Lempel , S. Moran  
**ACM Transactions on Information Systems (TOIS)** April 2001  
Volume 19 Issue 2  
Today, when searching for information on the WWW, one usually performs a query through a term-based search engine. These engines return, as the query's result, a list of Web pages whose contents matches the query. For broad-topic queries, such searches often result in a huge set of retrieved documents, many of which are irrelevant to the user. However, much information is contained in the link-structure of the WWW. Information such as which pages are linked to others can be used to augment search ...

---

Results 1 - 20 of 29    short listing

◀    ▶  
Prev    Next  
Page    Page  
1    2    \_

---

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.




[> home](#) [> about](#) [> feedback](#) [> login](#)

US Patent &amp; Trademark Office


 Try the *new* Portal design  
 Give us your opinion after using it.

## Search Results

Search Results for: [(updat\* or modify\* delet\* or insert\* or version\*)<AND>((report\*<AND>((scoreboard and electronic\* ) ) ) )]  
 Found 29 of 126,269 searched.

Search within Results

[> Advanced Search](#)[> Search Help/Tips](#)
 Sort by: [Title](#) [Publication](#) [Publication Date](#) [Score](#) [Binder](#)
Results 21 - 29 of 29 [short listing](#)

1

2




- 
- 21 [Exploiting data forwarding to reduce the power budget of VLIW embedded processors](#) 77%  
 M. Sami , D. Sciuto , C. Silvano , V. Zaccaria , R. Zafalon  
**Proceedings of the conference on Design, automation and test in Europe** March 2001
- 22 [A psychology of learning BASIC](#) 77%  
 Richard E. Mayer  
**Communications of the ACM** November 1979  
 Volume 22 Issue 11  
 This paper addresses the question: What does a person know following learning of BASIC programming? Several underlying conceptual structures are identified: (1) a transaction is an event that occurs in the computer and involves some operation on some object at some location, (2) a prestatement is a set of transactions corresponding to a line of code, (3) chunks are frequently occurring configurations of prestatements corresponding to several lines of code.
- 23 [A VLIW architecture for a trace scheduling compiler](#) 77%  
 Robert P. Colwell , Robert P. Nix , John J. O'Donnell , David B. Papworth , Paul K. Rodman  
**Proceedings of the second international conference on Architectural support for programming languages and operating systems** October 1987  
 Volume 15 , 22 , 21 Issue 5 , 10 , 4
- 24 [Pipeline Architecture](#) 77%  
 C. V. Ramamoorthy , H. F. Li  
**ACM Computing Surveys (CSUR)** January 1977  
 Volume 9 Issue 1
- 25 [The computer for the 21st century](#) 77%  
 Mark Weiser  
**ACM SIGMOBILE Mobile Computing and Communications Review** July 1999  
 Volume 3 Issue 3  
 Specialized elements of hardware and software, connected by wires, radio waves and infrared,

will be so ubiquitous that no one will notice their presence.

26 Reverse execution of programs

77%

 Bitan Biswas , R. Mall


**ACM SIGPLAN Notices** April 1999

Volume 34 Issue 4


Conventional debuggers do not allow users to go back and examine the program states at statements which have already been executed. In case the user wants to examine the program state at a statement which was executed sometime back, he is forced to restart the entire debugging process. To overcome this problem, we examine the issue of reverse execution of programs. To this end, we introduce the concept of *inverse of a statement*. We describe our implementation of a debugger which can execu ...

27 Creating a live broadcast from a virtual environment

77%

 Chris Greenhalgh , John Bowers , Graham Walker , John Wyver , Steve Benford , Ian Taylor  
**Proceedings of the 26th annual conference on Computer graphics and interactive techniques** July 1999

28 Exploiting fine-grained parallelism through a combination of hardware and software techniques 77%

 Stephen Melvin , Yale Patt

**ACM SIGARCH Computer Architecture News , Proceedings of the 18th annual international symposium on Computer architecture** April 1991

Volume 19 Issue 3

29 Liveboard: a large interactive display supporting group meetings, presentations, and remote collaboration

77%



 Scott Elrod , Richard Bruce , Rich Gold , David Goldberg , Frank Halasz , William Janssen , David Lee , Kim McCall , Elin Pedersen , Ken Pier , John Tang , Brent Welch

**Proceedings of the SIGCHI conference on Human factors in computing systems** June 1992

This paper describes the Liveboard, a large interactive display system. With nearly one million pixels and an accurate, multi-state, cordless pen, the Liveboard provides a basis for research on user interfaces for group meetings, presentations and remote collaboration. We describe the underlying hardware and software of the Liveboard, along with several software applications that have been developed. In describing the system, we point out the design rationale that was used to make various c ...

---

Results 21 - 29 of 29    short listing

   
Prev Page   1   2   Next Page

---

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.


[> home](#) [> about](#) [> feedback](#) [> login](#)

US Patent &amp; Trademark Office


 Try the *new* Portal design  
 Give us your opinion after using it.

## Search Results

Search Results for: [threshold\*<AND>(((updat\* or modify\* delet\* or insert\* or version\*)<AND>((report\*<AND>((scoreboard and electronic\* ) ) ) ) ) ]  
 Found 6 of 126,269 searched.

## Search within Results

[> Advanced Search](#)[> Search Help/Tips](#)
 Sort by: [Title](#) [Publication](#) [Publication Date](#) [Score](#) [Binder](#)
Results 1 - 6 of 6 [short listing](#)

- 1 [Talisman: fast and accurate multicomputer simulation](#) 89%  
 Robert C. Bedichek  
**ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1995 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems May 1995**  
 Volume 23 Issue 1  
 Talisman is a simulator that models the execution semantics and timing of a multicomputer. Talisman is unique in combining high semantic accuracy, high timing accuracy, portability, and good performance. This good performance allows users to run significant programs on large simulated multicomputers. The combination of high accuracy and good performance yields an ideal tool for evaluating architectural trade-offs. Talisman models the semantics of virtual memory, a circuit-switched interno ...
- 2 [Efficient dynamic scheduling through tag elimination](#) 84%  
 Dan Ernst , Todd Austin  
**ACM SIGARCH Computer Architecture News May 2002**  
 Volume 30 Issue 2  
 An increasingly large portion of scheduler latency is derived from the monolithic content addressable memory (CAM) arrays accessed during instruction wakeup. The performance of the scheduler can be improved by decreasing the number of tag comparisons necessary to schedule instructions. Using detailed simulation-based analyses, we find that most instructions enter the window with at least one of their input operands already available. By putting these instructions into specialized windows with fe ...
- 3 [Session 2: embedded system techniques \(1\): A case for dynamic pipeline scaling](#) 80%  
 Jinson Koppanalil , Prakash Ramrakhyani , Sameer Desai , Anu Vaidyanathan , Eric Rotenberg  
**Proceedings of the international conference on Compilers, architecture, and synthesis for embedded systems October 2002**  
 Energy consumption can be reduced by scaling down frequency when peak performance is not needed. A lower frequency permits slower circuits, and hence a lower supply voltage. Energy reduction comes from voltage reduction, a technique called Dynamic Voltage Scaling (DVS). This paper makes the case that the useful frequency range of DVS is limited because there is a lower bound on voltage. Lowering frequency permits voltage reduction until the

lowest voltage is reached. Beyond that point, lowering ...

- 4 Improving TCP performance over mobile networks 80%  
Hala Elaarag  
**ACM Computing Surveys (CSUR)** September 2002  
Volume 34 Issue 3  
*Transmission Control Protocol (TCP)* is the most commonly used transport protocol on the Internet. All indications assure that mobile computers and their wireless communication links will be an integral part of the future internetworks. In this paper, we present how regular TCP is well tuned to react to packet loss in wired networks. We then define mobility and the problems associated with it. We discuss why regular TCP is not suitable for mobile hosts and their wireless links by p ...
- 5 Illustrative risks to the public in the use of computer systems and related technology 80%  
Peter G. Neumann  
**ACM SIGSOFT Software Engineering Notes** January 1996  
Volume 21 Issue 1
- 6 Exploiting fine-grained parallelism through a combination of hardware and software techniques 77%  
Stephen Melvin , Yale Patt  
**ACM SIGARCH Computer Architecture News , Proceedings of the 18th annual international symposium on Computer architecture** April 1991  
Volume 19 Issue 3

---

Results 1 - 6 of 6    [short listing](#)

---

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.